Managing Responsibility and Information Flow in Dynamic Team Decision-Making

Paul Hubbard[*], **Alexander Kott and Michael Martin**
BBN Technologies
11 Stanwix St., Pittsburgh, PA 15222
412-690-2336 (2332)
{akott,mkmartin}@bbn.com

Abstract

We propose a method to allocate decision responsibility and arrange information flow dynamically within a team of decision-makers for command and control. A model of decision making which relates the decision load to decision accuracy is proposed and employed as an atomic building block to create predictive models of team decision making. An optimization problem is then proposed for a given set of forecast decision requirements in which the information flow between the atomic decision-making models is varied so as to maximize an aggregate measure of decision accuracy. A small-scale MATLAB Simulink implementation is presented as well as the outline of current work in which a genetic algorithm is employed to perform the optimization. Preliminary results indicate the technique improves the decision-making performance measure. We conclude with a discussion of implementation issues in a larger C2 context.

1. Introduction

We consider a scenario in which a distributed decision-making team such as a society of human warriors, robots and computer systems is engaged in execution of a common mission.   The decision-makers communicate with one another and adapt to changing circumstances, threats, opportunities and information requirements while attempting to achieve a common objective. They may also continually re-allocate decision-making tasks within the team and re-allocate information flow between decision-makers. This view differs from a more conventional, fixed hierarchical decision-making, or hierarchical control, and aims to account for the dynamic, fluid nature of decision responsibilities and interactions.

In the traditional theory and practice of distributed and hierarchical control systems, a supervising controller issues commands to lower level controllers. These commands become inputs or set points for the control processes of the lower-level controllers. The responsibilities and types of interactions between lower level controllers are static. They are fixed when the system is designed and remain constant while the system operates. When this constant definition of responsibilities and interactions becomes inadequate to adapt to the changing circumstances, as will be the case in an uncertain environment with adversarial threats and changing information

---

[*] Dr. Hubbard can now be reached at Defense R&D Canada – Ottawa, 3701 carling Ave., Ottawa, Ontario, Canada K1A 0Z4, (613)-730-5928, paul.j.hubbard@sympatico.ca.

Report Documentation Page		*Form Approved* *OMB No. 0704-0188*

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE **SEP 2002**	2. REPORT TYPE	3. DATES COVERED **00-00-2002 to 00-00-2002**

4. TITLE AND SUBTITLE **Managing Responsibility and Information Flow in Dynamic Team Decision-Making**	5a. CONTRACT NUMBER
	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)	5d. PROJECT NUMBER
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **BBN Technologies,11 Stanwix Street,Pittsburgh,PA,15222**	8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT **Approved for public release; distribution unlimited**

13. SUPPLEMENTARY NOTES **The original document contains color images.**

14. ABSTRACT

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT **unclassified**	b. ABSTRACT **unclassified**	c. THIS PAGE **unclassified**		**11**	

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std Z39-18

requirements, the control system is unable to perform effectively. Additionally, in the context of a team of decision-makers that is composed of humans as well as automated control processes, it is typically a human commander who controls the allocation of decision-making and information-sharing responsibilities as well as the modes of such sharing. The structure of the team is also often fixed in a hierarchical format. A human commander typically adjusts decision-making and information-sharing responsibilities based on informal techniques, intuition, and best guesses based on prior experience and training [Van Creveld, 1985], [Serfaty, 1998] [Entin,1999], [Carley, 1998]. However, this control can break down when the complexity and speed of changes in the situation exceeds the cognitive and reasoning capabilities of the human controller. This problem is exacerbated in those teams that include software agents and robots as decision-makers because these artificial decision-makers can observe, execute their decision-making algorithms, and act much faster than a human controller, leading to a cascade of failures and poor team performance [Morgan, 1995].

The above are among the examples of pathologies that can be exhibited by C2 systems. A catalog of such pathologies appeared in [Kott, 1999]. These range from clear and observable problems such as loss of synchronization or thrashing to more subtle problems in C2 processes such as hierarchical inconsistencies or positive feedback. An example of unstable behavior in a C2 system due to the positive feedback involving information overload was explored in [Kott, 2001]. The main feature of the pathology studied in [Kott, 2001] is that the C2 system, either internally or in interaction with the battlespace, enters into a self-reinforcing cycle of increasing decision workload until the demand for decision-making exceeds the capacity of the C2 system. The work also showed that by dynamically reallocating decision load between the decision-makers it is possible to expand significantly the envelope of stable operation. The same notion of a reduction in performance due to information overload is at play in the present work, though here we are looking at optimizing performance based on future decision requirements rather than developing a qualitative characterization of the operational envelope within which the decision process is stable.

Herein, we explore a computer-assisted approach for managing decision-making, information-sharing responsibilities, and modes of interactions between the team members that allows the team to effectively and rapidly adapt to changing circumstances, threats and opportunities.

2. **Predictive Control of Distributed Decision-Making**

The proposed predictive control scheme is illustrated in Figure 1 and involves four main aspects that are described in the following sections. In the setting of predictive control often employed, even if only implicitly, in planning and execution of military operations and industrial engineering applications, an internal model is used to determine, either analytically or through simulation or wargaming, an input that can then be applied to the real system. Here we consider the input to be the structure of the decision-making team, i.e. the decision responsibilities and information channels. The actual decision and control values are assumed to be internal to the physical battlespace as shown in Figure 1, i.e. we have abstracted away the epistemological content of the decisions themselves in order to capture quantitatively the decision load and information channel

loads. Similarly, as illustrated in Figure 1, the forecast decision requirements are abstracted to a set of time-varying parameters that are used as an open-loop input for the simulation of the decision-making process. The standard model-predictive control process is then enacted; different information structures are simulated with the given forecast decision requirements and a best structure is chosen and used for the real decision making team in the physical battlespace. This process of simulating a variety of information structures is captured by the internal optimization loop of Figure 1.
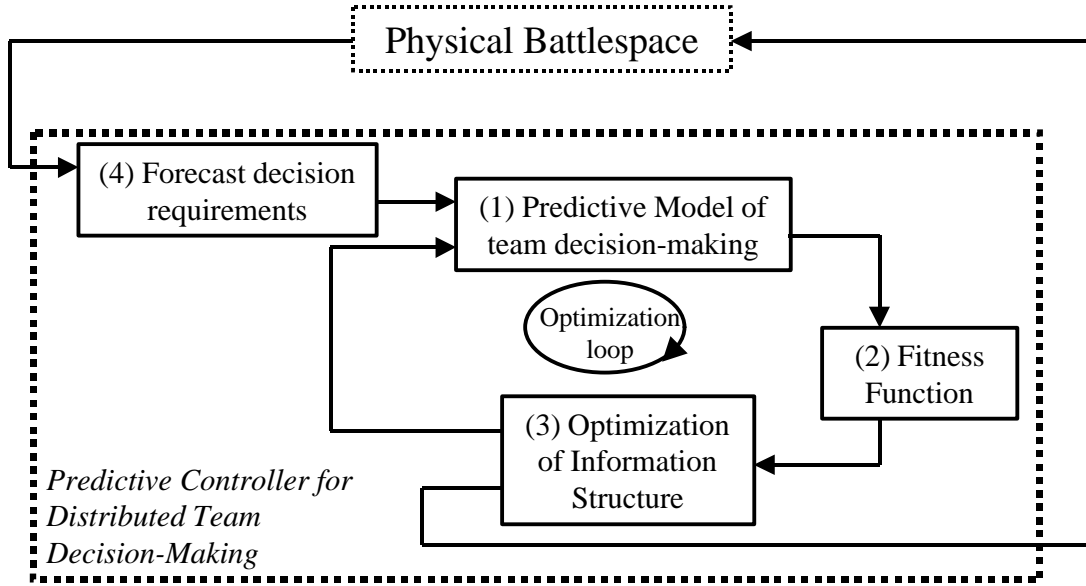

Figure 1. The proposed predictive control scheme for dynamic allocation of decision responsibilities and information sharing.

2.1 *An example: Three member team fighting a fire in a chemical plant.*

For illustrative purposes, we consider a simple example of a team of three decision-makers – a foreman, a scout and a robot – involved in a mission whose objective is to extinguish a fire in a large chemical plant. We assume that an initial ad-hoc plan calls for the foreman to observe the fire from an observation post, the scout to identify a target location for dropping fire extinguishing material and then join the foreman, and the robot to drop the fire extinguishing material at the identified target location. An example of possible decisions and observations required are as follows:

U1 – When and where to call a vehicle for the foreman escape.
U2 – When and on which target to drop fire extinguishing material.
U3 – Which flight path to take to reach the target.
U4 – When to egress and which egress route to choose for the scout.

We will re-visit this example as notions are developed. The questions we wish to address in the context of this example are

- Which, of the three team-members, should be responsible for deciding U1 through U4?
- How should the current decision outcomes be disseminated to other members in the team (noting, of course, the key fact that too much information has negative consequences when an urgent decision is needed)?
- If we can predict the difficulty and urgency of decisions U1 through U4, how should we manage dynamically the answers to the first two questions?

Note that the decision-makers in this example are one and the same as the actors, but this is for convenience only. There may be entities in the decision-making teams that are physically removed from the plant itself (an experienced foreman at another plant, for instance). Also note that we are not considering the control of the observation process[1], again for convenience.

3. **Modeling Team Decision-Making: Decision Responsibility and Information Structure**

At this investigative stage of the research, we employ a straightforward model to capture the traditional relationship between the task characteristics of complexity, urgency and decision load and the resulting accuracy of the decision [Luce, 1986], [Louvet, 1988], [Busemeyer, 1993], [Zsambok, 1997]. A set of parameterized curves in which accuracy is inversely proportional to the time pressure is used for this purpose. An example of this relationship is shown in Figure 2. The additional parameters of normalized discriminability and the number of the options provide an upper and lower bound on the accuracy as illustrated. In future work, we plan to investigate how the relationship between accuracy and decision requirements could be learnt or modeled analytically at a higher fidelity and study the relative impact of imperfect modeling of these relationships on the predictive power of the team decision-making model.

In a real implementation of the scheme, characteristics of the decision-making entities could be used to quickly generate models from a pre-determined virtual decision-maker coefficient database. Such characteristics might include rank and experience level for a human decision maker or processor capability, function and speed for an artificial decision-maker. As we are presently interested in the analysis of trends and the effectiveness of the approach, a relatively simple model suffices.

The relationship in Figure 2 is used to build dynamic input/output models for single decision-makers. Currently, a discrete-time approach is used; the accuracy, at a particular time instant, of outgoing decisions for which that decision-maker is responsible is a function of the decision requirements at the previous time instant. The normalized discriminability and number of options are exogenous inputs to the decision-making team and we assume these are properties of the decisions themselves rather than the interaction processes in the team decision-making. Further, we assume the time-pressure has two sources. First, there is an exogenous component that must be predicted, but additionally there is time-pressure that results from the necessity of interactions with other team members. Clearly, if another team-member is urgently soliciting information on a

---

[1] In a patent application on the same topic [Kott, 2002], it is proposed that observation channels be controlled in much the same way as decisions. In this example, this could mean appending the observations 'Z1- Location of the fire' and 'Z2 – Types of target locations' for instance to the list of decisions.

different topic, this reduces the amount of time that can be spent on the decision at hand and so increases the time pressure.
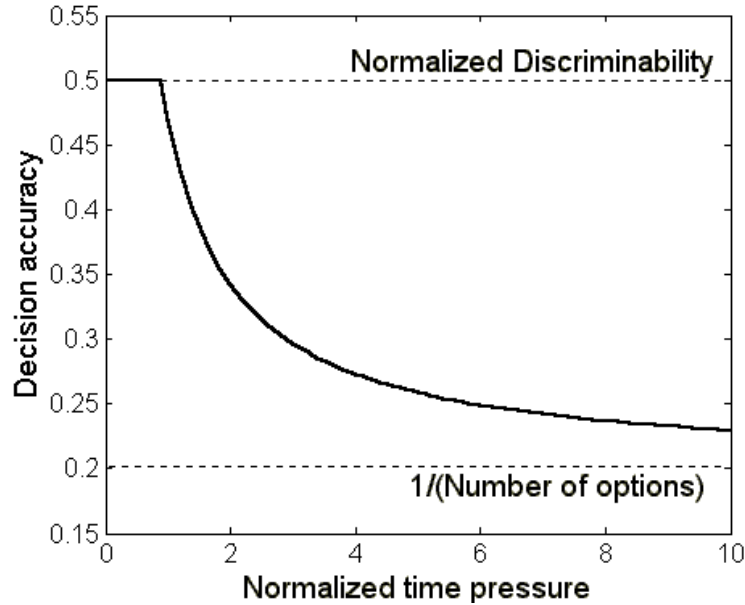

Figure 2. An example of the decision-accuracy relationship for one decision-maker.

The input/output model for a single decision-maker is used as an atomic block from which is constructed the team decision-making model. The connections between the atomic blocks are determined by what we will term the *information structure*. This term is not uncommon in the literature on decentralized and distributed control in more general settings [Witsenhausen, 1971], [Teneketzis, 1997], [Ho, 1980]. We borrow the term here to refer specifically to the combination of three objects: an *observation structure*, which maps the incoming data streams to the decision-making entities, or agents, a *decision-responsibility structure* which partitions the outgoing decision variables amongst the agents and a *decision-sharing structure* which defines the information paths between agents.

A possible information structure for the fire-fighting example discussed above is partially illustrated in Figure 3. The decisions that each member is responsible for are indicated on the figure and this illustrates the decision responsibility structure. The decision-sharing structure is illustrated with connecting arcs. The labels on these arcs show the decision that is being communicated and the *mode* in which it is communicated. Here, we consider only two modes: a *pull* mode in which the first decision-maker communicates information only if the second requests the communication, and a *push* mode, in which the first decision-maker immediately communicates any new information. Though there is certainly a spectrum of possible interactions of this sort, we have identified these two modes for a preliminary analysis. The observation structure is again left out for convenience (see the previous footnote). Note, from a semantic point of view, there is a minimal necessary information transfer for the decisions to be made. For instance, decision U3 (the route taken to the target location) necessitates knowledge of decision U2 (selection of the location and timing for the target). Accordingly, we define a *decision dependency criterion* that underlies the dynamics of the model. Information structures that do not

meet this minimal criteria will perform poorly because, for example, decisions concerning U3 made in the absence of knowledge of U2 are given arbitrarily an accuracy of zero per cent (we assume informally that two wrongs do not make a right). The decision dependencies for this example are that U1 requires knowledge of U2 and U3, U2 requires U1, U3 requires U1 and U2 and U4 requires U1 and U2. We assume that the minimal criterion can be met with either the pull or push mode communication.
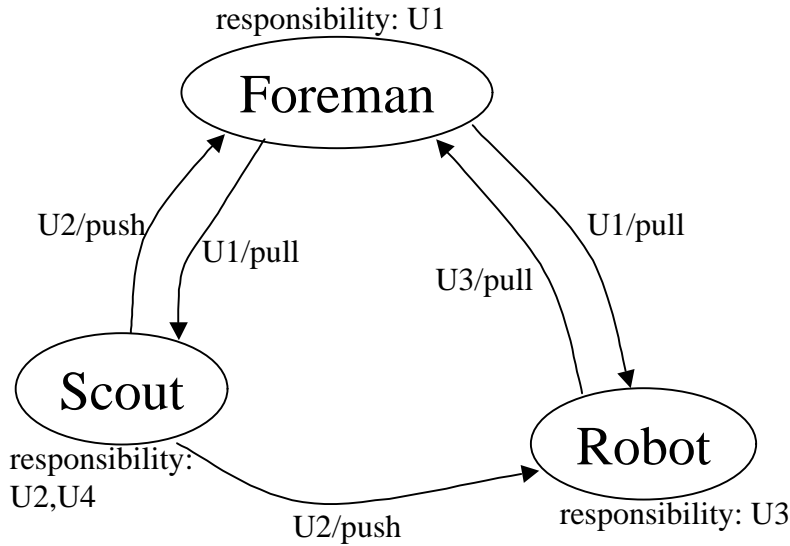

Figure 3. The decision-makers in the fire-fighting example communicate with one another and take responsibility for specific tasks as described by an information structure.

4. **Fitness Function**

A quantitative measure of decision-making performance (e.g. overall accuracy or timeliness of decisions) is required as a measuring stick for the suitability of a given information structure. Clearly, the true intrinsic value of a given information structure is its effectiveness in the battlespace or for the task at hand, but unfortunately without a high-fidelity battlespace simulation this intrinsic value is not available in a simulation of the decision-making team. While a number of surrogate measures can be envisioned, we explore a particular one - a weighted average of the accuracies of all decisions being made by the team, integrated over the time horizon, that serves as an approximate indication of performance. A potential implementation would couple decision-making models with a battlespace simulation so that the predicted overall effectiveness could in fact be used as a fitness function.

As indicated in Figure 1, the value of the fitness function is used to fine-tune the information structure. Smoothness or monotonicity of the fitness function with respect to changes in information structure will not generally hold. More importantly, changes in the information structure are discrete yielding a discrete (in state) optimization problem and highlighting the fact that traditional continuous-state feedback control techniques are not directly applicable.

5. **Optimization of Information Structure and the Optimization Loop**

Repeated simulation for different alternative information structures is performed to maximize the fitness function. This in effect requires a global search of the space of information structures. The optimization problem can be stated formally in a form similar to the observation problem suggested in [Ho,1980] and there, as here, it bears no analytic solution. In practice, the optimization of responsibility/information structure cannot be de-coupled from the optimization of the physical actions, such as the execution of a military operation. While traditional operations planning tools do not explicitly consider the information structure as a quantity to be controlled, a two-step process could be employed to optimize physical operations serially, either before or after, the optimization of information structure.

6. **Forecast Decision Requirements**

In our current prototype implementation described below, the forecast decision characteristics consist of a normalized discriminability, the number of options and a time pressure that characterize each decision. These are predicted for a rolling time horizon (for the next 6 hours, for instance) and these predictions are then used as open-loop inputs to the dynamic model of decision-making. The feedback that would lead to changing predictions due to battlespace events that are a result of the change in information structure is not considered in the current work. The determination of which decision characteristics are of most importance for predicting decision performance is also beyond the scope of the current work though of and great interest and vital importance for the success of the proposed approach.
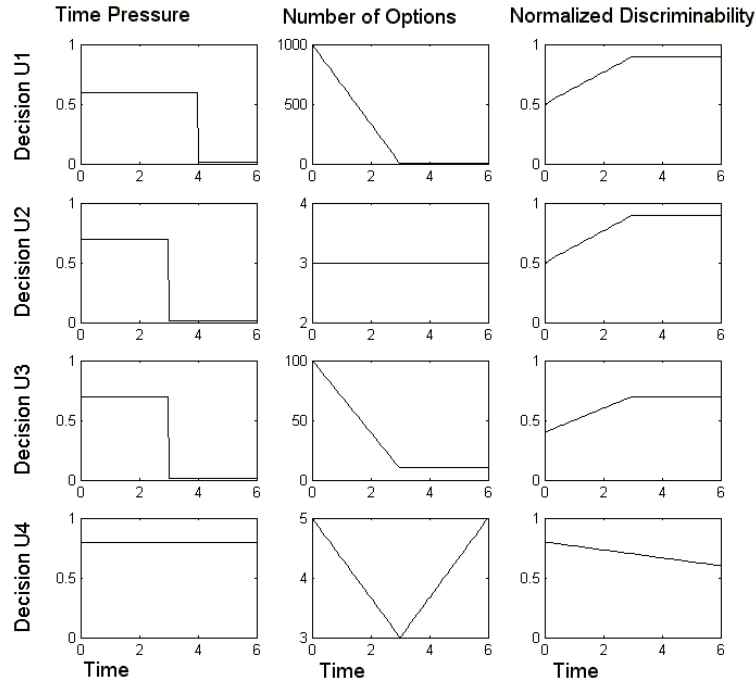

Figure 4. Future decision characteristics are predicted from information sources and used to fine tune the information structure.

In practice, the raw information from which these parameters could be predicted will come directly from the observations of the battlespace and other information and intelligence sources.

An example of predicted decision requirements for the fire-fighting example are shown in Figure 4 for a time horizon of 6 hours. For the purposes of our example, these were generated manually on the basis of the notional progression of the fire and the related actions of the fire-fighting team. For instance, the decision characteristics of U1 (corresponding to the escape vehicle decisions for the foreman) are such that there is little time pressure and many options at time 0, but as the fire progresses the time pressure increases (reflected by a decrease in the graph) and the number of options decreases steadily. This might be the case if escape vehicle options become scarcer as the fire approaches locations close to the foreman.

7. A Prototype Implementation

Figure 5 shows a Matlab Simulink [MathWorks, 2002] prototype implementation of the model in which atomic blocks for the individual decision-makers in the fire-fighting example are combined with the information structure in Figure 3.
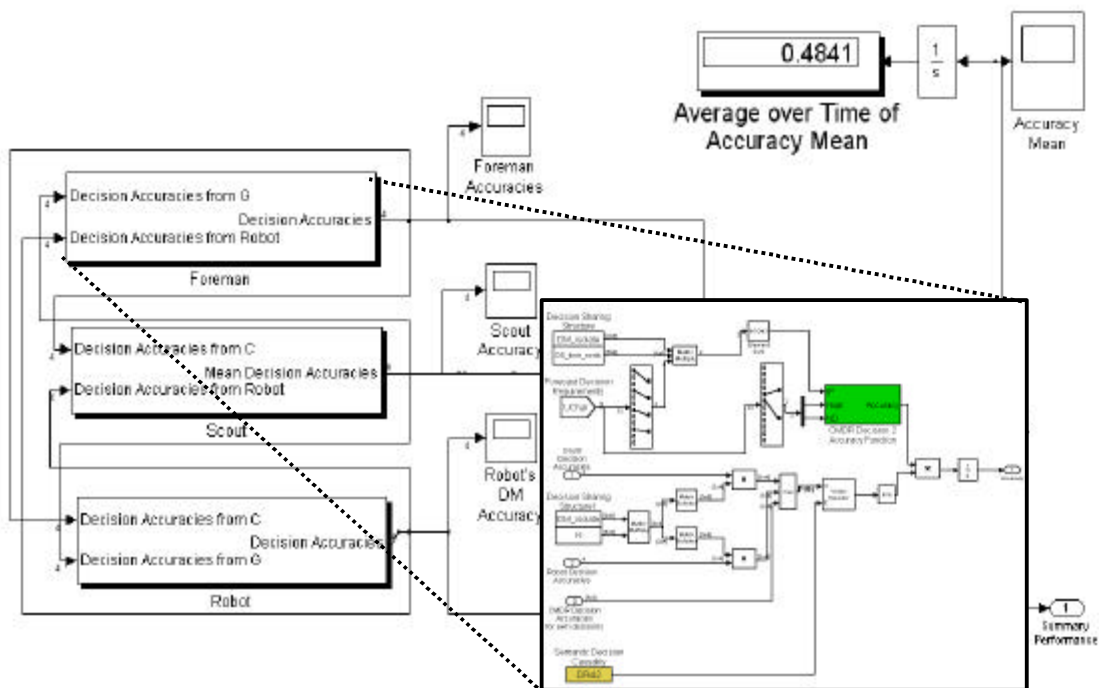

Figure 5. The proposed scheme for predictive control implemented
in MATLAB Simulink for a small-scale example.

We are currently investigating an optimization scheme based on a genetic algorithm although a number of discrete optimization techniques could be potentially applicable. The information structure is encoded as a chromosome without content loss, and the fitness of each chromosome is evaluated with a direct simulation of our team-decision making MATLAB Simulink model. The Genetic Algorithm Optimization Toolbox [GAOT Toolbox, 2002] offers a large variety of mutation operators, crossover operators and selection criteria and is well suited for this application.

Preliminary results for this optimization appear promising for the small-scale exploratory scenario with convergence to a (locally) optimal information structure with greatly improved fitness in realistic computational cost (less than one minute on a Pentium 4 Processor). A small variety of meaningful, common-sense, information structures were produced (see Figure 6 for examples), depending on how the forecast predictions were varied, indicating a match at least with intuitive notions of how teams should communicate and share tasks.
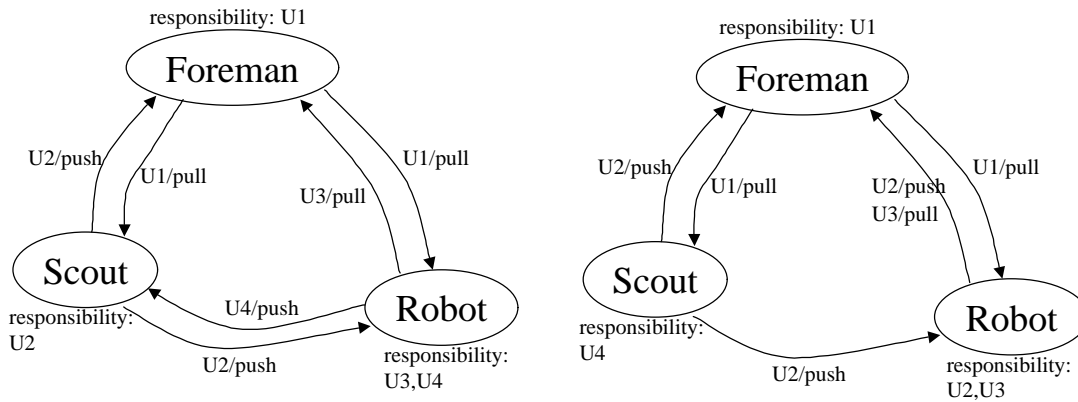

Figure 6. Examples of locally optimal information structures
for given sets of predicted decision characteristics.

.

8. **Conclusions**

We argue that dynamic modification of the decision responsibilities and information-sharing links within a decision-making team can improve stability and performance in terms of quality of decisions produced by the team. We have proposed a model-predictive scheme to control the modifications during operations based on (continually updated) forecast decision complexities and urgencies. The scheme involves modeling individual decision-makers and combining them to produce a predictive model of the team information exchange and decision-making. Figure 7 shows an informal landscape of the applicability of this form of dynamically managed team decision-making. Intuitively, as the tempo of operations increases and the accuracy of predictions in the battlespace decrease, it is expected that dynamic management of the information structure

will improve performance. In particular, the proposed approach would be suitable to teams composed entirely or partially of unmanned or semi-autonomous vehicles.

We have implemented and experimented with small-scale prototypes of team decision-making models. Though these have provided favorable results, the prototypes were implemented without regards to the impact of issues such as modeling uncertainty, inaccuracies in predicted decision requirements and possible counter-intuitive dynamics evolving from human-computer interactions. It is clear that these and other issues need to be addressed in future research.
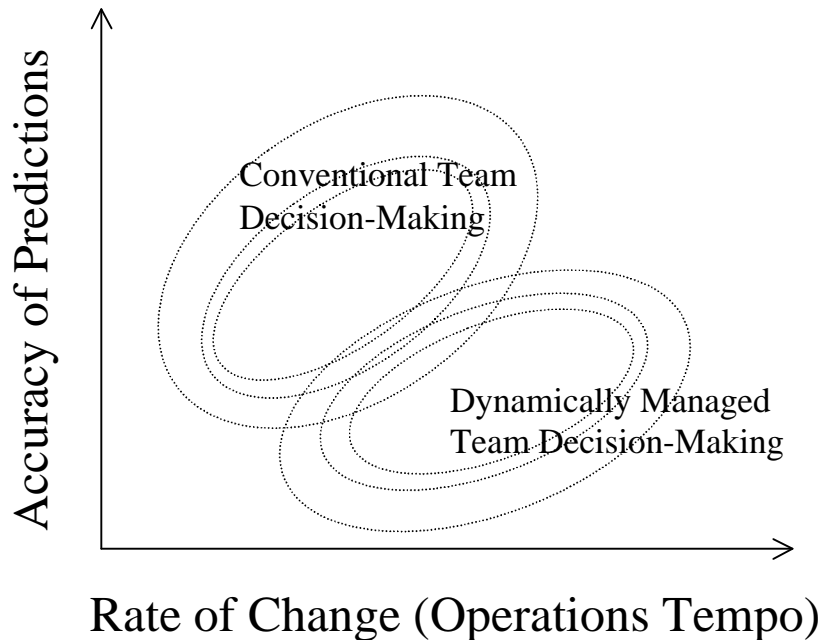

Figure 7. The regions of general applicability of conventional and dynamically managed team decision-making.

8. **References**

[Kott, 2002] Kott A., Hubbard, P., and Martin, M., Method and system for dynamic allocation of roles and information sharing within a decision-making team, Patent Application, BBN Technologies, 2002

[Kott, 1999] Kott, A. and Krogh, B., Toward a Catalog of Pathological Behaviors in a Complex Enterprise Control System, Proceedings of the DARPA-JFACC Symposium on Advances in Enterprise Control, San Diego, November 1999.

[Kott, 2001] Kott, A., Hubbard, P. and Martin, M., Instability of Distributed Decision Making in Command and Control Systems, Proc. of 2001 American Control Conference, June 2001.

[Morgan, 1995] Morgan, B.B. Jr., & Bowers, C.A.. Teamwork stress: Implications for team decision making. In R.A. Guzzo, E. Salas, and Associates (Eds.) Team effectiveness and decision making in organizations. San Francisco: Jossey-Bass Publishers, 1995.

[Van Creveld, 1985] Van Creveld, M., Command in War, Harvard University Press, Cambridge, MA,1985.

[Serfaty, 1998] Serfaty, D., Entin, E. E., & Johnston, J. H. Team coordination training. In J. A. Cannon-Bowers & E. Salas (Eds.), Making decisions under stress: Implications for individual and team training (pp. 221-245). Washington, D. C.: American Psychological Association, 1998.

[Entin,1999] Entin, E.E.. Optimized Command and Control Architectures for Improved Process and Performance. Proceedings of the 1999 Command and Control Research and Technology Symposium, Newport, RI., 1999.

[Carley, 1998] Carley, K.M., Prietula, M.J. & Zhiang, L., Design versus cognition: the interaction of agent cognition and organizational design on organizational performance, Journal of Artificial Societies and Social Simulation, 1(3):1-19 June 1998.

[Luce, 1986] Luce, R.D.. Response times: Their role in inferring elementary mental organization. New York: Oxford University Press, 1986.

[Louvet, 1988] Louvet, A-C., Casey, J. T., and Levis, A. H, Experimental Investigation of the Bounded Rationality Constraint, in S. E. Johnson and A. H. Levis (ed.), Science of Command and Control: Coping with Uncertainty, AFCEA, Washington, D.C., pp. 73-82, 1988.

[Busemeyer, 1993] J.R. Busemeyer & J.T. Townsend, Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. Psychological Review, 100, 432-459, 1993.

[Zsambok, 1997] C.E. Zsambok, G. Klein, Ed.,   Naturalistic Decision Making, Laurence Erlbaum Associates, NJ, 1997.

[MathWorks, 2002] MathWorks Website: http://www.mathworks.com/, 2002.

[Witsenhausen, 1971] H.S. Witsenhausen. On Information Structures, feedback and causality, SIAM J. Control , 9 2, May 1971.

[Teneketzis, 1997] D. Teneketzis, On Information Structures and Nonsequential stochastic control, CWI quarterly, 10 (2), 1997.

[Ho, 1980] Y-C Ho, Team Decision Theory and Information Structures, Proceedings of the IEEE, 62 6, June 1980.

[GAOT Toolbox, 2002] Genetic Algorithm Optimization Toolbox (GAOT) for Matlab 5: http://www.ie.ncsu.edu/mirage/GAToolBox/gaot/gaotindex.html, 2002.